

Ask Bob Kahn:

MORE ON THE MEMOS THAT YOU HAD GENERATED ON N'WORKING WHILE AT BBN...AND THE LETTER YOU DRAFTED TO ROBERTS....

From Babbage Interview:

He certainly had my papers. In fact, as I was writing stuff I generally would send it to him. I can even recall sending him one paper that basically said whatever he did with the network, he ought to have a long line involved in it up front so he could find out whether the network would work with long distance lines or not. Because everybody was thinking about small scale experimentation. I was afraid if they only did a small scale experiment that it might work in the small but fail in the large. So I felt they ought to set the initial structuring of it to be consistent with what the eventual network would look like. Worry about the network on the scale that you really want to operate it. You know, if you are going to deal with real errors make sure you are going to encounter approximately the same number you would expect in a reasonable number of configurations. It seemed to me at the time that the longest lines you would ever want in a network would be roughly a few thousand miles long. So put in one or more of them initially, just to see how it would work.

O'NEILL: You came into networking from a theoretical point of view, is that right?

KAHN: I would say from a conceptual point of view, yes. I mean, there was very little theory at the time.

MORE ON YOUR WORK ON ROUTING ALGORITHMS?

MORE ON THE SIMULATION PROGRAM?

But I think what finally triggered my going over there was an activity that involved simulation. Along with another colleague at BBN, we developed a simulation program that could be used for evaluating network performance. Most of the really interesting issues to me about networks were how they were going to work, because this was a really new area. Nobody had any experience in knowing how networks would actually function in practice: how did routing algorithms work, how did congestion control algorithms work, how did flow control work, how did buffering work. All of those kinds of things were critical issues back then. The simulation system was developed to give us a visual clue. I mean, this was a rather innovative development because - remember this was at time when interactive graphics were not par for the course. I mean, there were no such things as workstations. And time-sharing was pretty new at the time as

well. So to be able to get an interactive environment with graphics displays of networks was a major coup in its own right. So we had a very powerful facility to be able to use, and my thought was to use it to explore some of the issues.

MORE ON YOUR MEMORY OF THE WRITING OF 1822?

MORE ON THIS?

And, in fact, one of my reasons for accepting the DARPA job was because it just didn't look like BBN was interested in capitalizing on any of the technology we had developed.

One of the things that I chose not to do was to run any of the network project. Having just come from BBN, I wanted to get into new things. In fact, the agreement that I had when I went to DARPA was I would set up a program in flexible manufacturing. I was intending to make a clean break from networking. Apparently quite a bit of money had been set aside, or planned, or budgeted, for the new program, but when I got there it had been canceled by the Congress. So the program somehow disappeared in real-time right in front of my eyes. I remember Larry coming to me and saying "Look, I know you didn't want to work on networking anymore, but you know more about it than anybody else around and that's where our main efforts are going to be for the next several years, so why don't you just go do that." So I did, but I got into all the new efforts.

MORE ON THE DEADLOCK TESTS W/WALDEN? LATE '69, EARLY '70?

MORE ON THE MULTI-LAYER PROTOCOLS YOU WERE TALKING ABOUT?

YOUR MEMORY OF THE BACK-OF-THE-ENVELOPE STUFF W/VINT?

PACKET RADIO VS. ALOHA? PARALLEL EFFORTS?

The idea for a radio-based version of ARPANET, and a satellite-based version, had also been floating around the ARPA office for a while. I think there may even have been a small study contract given to one or two organizations to see if it was feasible, before I got to DARPA. When I got there there was money budgeted for a packet radio program, and I undertook to make it happen. The skids were all greased for that. Part way through the first year of the program it became clear to me that we were going to have to have a plan for getting computer resources on the net. In 1973, mainframe computers were multi-million dollar machines that required air-conditioned computer centers. You weren't going to connect them to a mobile, portable packet radio unit and carry it around.

So my first question was "How am I going to link this packet radio system to any computational resources of interest?" Well, my answer was, "Let's link it to the ARPANET." Except that these were two radically different networks in many ways. I mean, all the details were different. I don't mean conceptually they were different. They were sort of the same genre. Just like, say, Chinese and Americans are of the same genre except one speaks Chinese and one speaks English, one lives on one side of the world, one lives on the other side, they go to sleep during your daytime, etc. The details of the two

networks were rather different. The ARPANET ran at 50 kilobits per second and the packet radio system ran at 100 or 400 kilobits per second. One had thousand bit uncoded packets; the other had two thousand bit packets which could be coded. The ARPANET assumed that once you sent something it was delivered with a hundred percent reliability. If it didn't get through the system was broken. The other assumed that much of the time you would never get anything through even though the system was working. The protocols that were designed for the ARPANET wouldn't work over the packet radio net because when a packet entered the packet radio net, the only thing the ARPANET would have told it was where it came from but not where it was going. So the packet radio net had no further information to know where to route it. If a packet got lost along the way, the ARPANET hosts would come to a halt. Well, in a radio net you can get interference and some loss is natural. So we really had to rethink literally the whole issue of host transport protocols. Vint Cerf and I jointly came up with the TCP/IP concept as a new transport mechanism as part of an architecture for internetworking. DARPA then gave a contract to Vint at Stanford to actually implement the TCP/IP concept - along with small efforts at BBN and at University College London. Vint had the lead for developing the specification.

O'NEILL: WHAT WERE THE OTHER TWO CONTRACTS FOR?

KAHN: The one at University College London was part of an effort over there to explore remote use of the ARPANET and to implement a TCP/IP protocol on their machine. I think it was a PDP9 at the time. BBN was under contract to build a piece of the packet radio system (called a station) that needed portions of this protocol. That was the first embodiment of the notion of a gateway. We needed to implement a protocol that would work across the gateway. Eventually we all took the Internet technology pieces and created a separate program in DARPA for it. But originally, all that work was done as part of the packet radio program. I subsequently hired Vint Cerf to come to DARPA and actually run the Internetting program. That was in 1976, and by that time we were already three years into it. So he took it from a fledgling effort and turned it into a major national activity.

MORE ON THIS:

O'NEILL: Had you discussed it with other people in the community as well?

KAHN: Not really. Just the two of us. On occasion, we bounced ideas off of other people, but it was pretty much just Vint and myself working in a little room at what was then the Cabana Hyatt in Palo Alto. He had also come to Washington a few times. We met at Dulles Airport and did some work there, and I had gone out to Palo Alto. So we had actually had a few sessions together and then at one point during the summer of 1973 we just sat ourselves down for an entire weekend and just wrote the paper.

Ask Bob Kahn:

MORE ON THE MEMOS THAT YOU HAD GENERATED ON N'WORKING WHILE AT BBN...AND THE LETTER YOU DRAFTED TO ROBERTS....

From Babbage Interview:

He certainly had my papers. In fact, as I was writing stuff I generally would send it to him. I can even recall sending him one paper that basically said whatever he did with the network, he ought to have a long line involved in it up front so he could find out whether the network would work with long distance lines or not. Because everybody was thinking about small scale experimentation. I was afraid if they only did a small scale experiment that it might work in the small but fail in the large. So I felt they ought to set the initial structuring of it to be consistent with what the eventual network would look like. Worry about the network on the scale that you really want to operate it. You know, if you are going to deal with real errors make sure you are going to encounter approximately the same number you would expect in a reasonable number of configurations. It seemed to me at the time that the longest lines you would ever want in a network would be roughly a few thousand miles long. So put in one or more of them initially, just to see how it would work.

O'NEILL: You came into networking from a theoretical point of view, is that right?

KAHN: I would say from a conceptual point of view, yes. I mean, there was very little theory at the time.

MORE ON YOUR WORK ON ROUTING ALGORITHMS?

MORE ON THE SIMULATION PROGRAM?

But I think what finally triggered my going over there was an activity that involved simulation. Along with another colleague at BBN, we developed a simulation program that could be used for evaluating network performance. Most of the really interesting issues to me about networks were how they were going to work, because this was a really new area. Nobody had any experience in knowing how networks would actually function in practice: how did routing algorithms work, how did congestion control algorithms work, how did flow control work, how did buffering work. All of those kinds of things were critical issues back then. The simulation system was developed to give us a visual clue. I mean, this was a rather innovative development because - remember this was at time when interactive graphics were not par for the course. I mean, there were no such things as workstations. And time-sharing was pretty new at the time as

well. So to be able to get an interactive environment with graphics displays of networks was a major coup in its own right. So we had a very powerful facility to be able to use, and my thought was to use it to explore some of the issues.

MORE ON YOUR MEMORY OF THE WRITING OF 1822?

MORE ON THIS?

And, in fact, one of my reasons for accepting the DARPA job was because it just didn't look like BBN was interested in capitalizing on any of the technology we had developed.

One of the things that I chose not to do was to run any of the network project. Having just come from BBN, I wanted to get into new things. In fact, the agreement that I had when I went to DARPA was I would set up a program in flexible manufacturing. I was intending to make a clean break from networking. Apparently quite a bit of money had been set aside, or planned, or budgeted, for the new program, but when I got there it had been canceled by the Congress. So the program somehow disappeared in real-time right in front of my eyes. I remember Larry coming to me and saying "Look, I know you didn't want to work on networking anymore, but you know more about it than anybody else around and that's where our main efforts are going to be for the next several years, so why don't you just go do that." So I did, but I got into all the new efforts.

MORE ON THE DEADLOCK TESTS W/WALDEN? LATE '69, EARLY '70?

MORE ON THE MULTI-LAYER PROTOCOLS YOU WERE TALKING ABOUT?

YOUR MEMORY OF THE BACK-OF-THE-ENVELOPE STUFF W/VINT?

PACKET RADIO VS. ALOHA? PARALLEL EFFORTS?

The idea for a radio-based version of ARPANET, and a satellite-based version, had also been floating around the ARPA office for a while. I think there may even have been a small study contract given to one or two organizations to see if it was feasible, before I got to DARPA. When I got there there was money budgeted for a packet radio program, and I undertook to make it happen. The skids were all greased for that. Part way through the first year of the program it became clear to me that we were going to have to have a plan for getting computer resources on the net. In 1973, mainframe computers were multi-million dollar machines that required air-conditioned computer centers. You weren't going to connect them to a mobile, portable packet radio unit and carry it around.

So my first question was "How am I going to link this packet radio system to any computational resources of interest?" Well, my answer was, "Let's link it to the ARPANET." Except that these were two radically different networks in many ways. I mean, all the details were different. I don't mean conceptually they were different. They were sort of the same genre. Just like, say, Chinese and Americans are of the same genre except one speaks Chinese and one speaks English, one lives on one side of the world, one lives on the other side, they go to sleep during your daytime, etc. The details of the two

networks were rather different. The ARPANET ran at 50 kilobits per second and the packet radio system ran at 100 or 400 kilobits per second. One had thousand bit uncoded packets; the other had two thousand bit packets which could be coded. The ARPANET assumed that once you sent something it was delivered with a hundred percent reliability. If it didn't get through the system was broken. The other assumed that much of the time you would never get anything through even though the system was working. The protocols that were designed for the ARPANET wouldn't work over the packet radio net because when a packet entered the packet radio net, the only thing the ARPANET would have told it was where it came from but not where it was going. So the packet radio net had no further information to know where to route it. If a packet got lost along the way, the ARPANET hosts would come to a halt. Well, in a radio net you can get interference and some loss is natural. So we really had to rethink literally the whole issue of host transport protocols. Vint Cerf and I jointly came up with the TCP/IP concept as a new transport mechanism as part of an architecture for internetworking. DARPA then gave a contract to Vint at Stanford to actually implement the TCP/IP concept - along with small efforts at BBN and at University College London. Vint had the lead for developing the specification.

O'NEILL: WHAT WERE THE OTHER TWO CONTRACTS FOR?

KAHN: The one at University College London was part of an effort over there to explore remote use of the ARPANET and to implement a TCP/IP protocol on their machine. I think it was a PDP9 at the time. BBN was under contract to build a piece of the packet radio system (called a station) that needed portions of this protocol. That was the first embodiment of the notion of a gateway. We needed to implement a protocol that would work across the gateway. Eventually we all took the Internet technology pieces and created a separate program in DARPA for it. But originally, all that work was done as part of the packet radio program. I subsequently hired Vint Cerf to come to DARPA and actually run the Internetting program. That was in 1976, and by that time we were already three years into it. So he took it from a fledgling effort and turned it into a major national activity.

MORE ON THIS:

O'NEILL: Had you discussed it with other people in the community as well?

KAHN: Not really. Just the two of us. On occasion, we bounced ideas off of other people, but it was pretty much just Vint and myself working in a little room at what was then the Cabana Hyatt in Palo Alto. He had also come to Washington a few times. We met at Dulles Airport and did some work there, and I had gone out to Palo Alto. So we had actually had a few sessions together and then at one point during the summer of 1973 we just sat ourselves down for an entire weekend and just wrote the paper.